

2003

Insect Small-Target Motion Detection for Seeker Applications

Final Report

Contract N00014-03-M-0171

**Patrick Shoemaker
David O'Carroll**

Tanner Research, Inc.
2650 East Foothill Blvd.
Pasadena, CA 91107

This report is copyright 2003.

DISTRIBUTION STATEMENT A. APPROVED FOR PUBLIC RELEASE; DISTRIBUTION IS UNLIMITED.

20031117 043

1 Identification and Significance of the Problem or Opportunity

Although insects are relatively simple organisms compared to vertebrates, with nervous systems of limited size and complexity, they nonetheless possess capabilities that could greatly enhance the performance of autonomous flying vehicles and weapons if they could be duplicated in an artificial system. Insects do a remarkable job of controlling flight and other behaviors based on their low-resolution visual sense. Computation of optical flow for estimation of egomotion, and detection and tracking of moving targets, are two examples of such processing. For man-made systems that emulate these capabilities, we would expect to find applications in guidance and control and in seeker technology.

Wide-field neurons that respond to broad patterns of optical flow have been the most widely studied motion-sensitive cells in the insect nervous system to date, but neurons that respond preferentially to isolated moving targets have also been described. The best-understood example of these are the 'figure detecting' (FD) cells in dipterans [7][22], which are sensitive to moving objects up to tens of degrees in extent, and which are thought to play a role in detecting parallax motion of relatively nearby objects with respect to ground. However, neurons that respond selectively to very small moving targets have also been found in insects of several species, which typically pursue prey or mates as part of their normal behavior [19]. These neurons have been labeled Small Target Movement Detectors (STMDs). Their selectivity for small targets is particularly significant because, given the relatively low resolution of insect vision, moving objects of interest (prey, mates, or rivals) essentially remain point targets until they are in near proximity to the animal observing them. This characteristic might also serve as an object lesson for the development of artificial seekers: while the historic drive has been toward higher spatial resolution in the sensor, with the assumption that higher performance is a natural outgrowth of the massive amounts of raw data, the insect shows that smart processing in conjunction with a low-resolution sensor is capable of remarkable performance. Size and power consumption of the sensor is naturally a significant factor with regard to the miniaturization (and cost) of autonomous vehicles and weapons.

The STMD neurons were first described within the last decade, and as yet are not as well understood as the FD system. However, they are currently the subject of intensive study by the co-investigators on this contract, under another DoD-funded research project (US Air Force Office of Scientific Research contract F49620-01-C-0030). This Air Force contract has supported basic research on the STMD cells, and initial efforts to model them. It has also led to a significant discovery, made since the commencement of the current SBIR contract: the ability of at least some STMD neurons to respond to small moving targets in the presence of moving cluttered background, while rejecting background motion alone. Because this remarkable 'moving target / moving background' capability is so clearly applicable to autonomous weapons that must acquire and track targets while in near-ground flight, understanding and modeling it is of paramount interest.

In this project, we proposed to study and model STMD neurons as applied to imagery relevant to the problem of air-to-surface tracking of a moving target. This contract leverages heavily off of the Air Force project mentioned above; it has presented an opportunity for the development of new data handling and simulation tools, investigation of scenarios relevant to the

air-ground problem, and it has supported initial efforts to model the 'moving target / moving background' capability. This work is seen as a first step toward biomimetic seeker technology that will eventually enhance autonomy and performance of platforms such as 2.75-inch rockets, in a program such as LOGIR.

With the opportunity afforded by this research are associated risks, and we identified the novelty of the subject as the most significant of these risks: this is an SBIR project that focuses on a topic with some aspects that are still legitimate objects of basic research. This leaves a chance that some important, unknown aspect of STMD function may not be entirely resolved within the scope of follow-on efforts (although we are convinced that our Phase I effort will serve to lessen such doubts). With respect to the opportunity presented, we remain convinced that it is impossible to examine the known physiology of STMD cells without being convinced that they play an important role in just the sort of capabilities that are crucial for autonomous flying vehicles and weapons.

2 Phase I Technical Objectives

The original technical objectives of the proposed effort are as follows:

1. Examine and characterize the responses of STMDs to moving imagery relevant to the air-to-surface problem, i.e., small targets moving within a diverging optical flow field. (This objective has been modified to include targets moving against simpler, more uniform optic flow fields, and actual air-ground imagery of moving targets.)
2. By modeling STMDs and considering their interface with both front end and higher-level processing, develop a scenario for applying them to the air-to-surface problem.

3 Phase I Accomplishments

Three tasks were specified under the Phase I work plan: a Neurobiology task, an STMD Modeling task, and a Higher-Level Issues task. Accomplishments are discussed below according to each task.

3.1 Neurobiology

Neurobiological work included *in-vivo* intracellular recordings made from putative STMD neurons in intact animals, in response to displayed moving imagery. In this procedure, the animals are immobilized, a small portion of the rear head capsule is removed, and a drawn glass microelectrode filled with electrolyte solution, with a tip of less than 100nm diameter, is inserted to penetrate individual neurons. The rear head is bathed in physiological saline and a reference electrode placed in this bath. The membrane potential is recorded as moving visual scenarios are displayed to the animal on a CRT at a 200 frame per second rate.

The video imagery consisted of moving targets in artificially generated scenarios, and also natural imagery of moving targets from an air-to-ground perspective, supplied by NAWC early on in the project. The actual experiments with artificial imagery were conducted under support of Air Force contract F49620-01-C-0030, whereas analysis of the data obtained from them was

supported by the present contract. Particular attention was paid to the responses of STMD neurons to moving targets against moving backgrounds.

Progress in experimental biology was significantly hampered by a late start of work under this contract. Whereas we anticipated a mid-summer start (mid-January to early February in Australia where the O'Carroll laboratory is located), contract award took place at the end of March. Availability of one of the subject species (the dragonfly *Hemicordulia tau*) rapidly diminishes during the austral autumn, and none could be obtained during the period of performance. However, we did succeed in finding a small, irregular supply of hoverfly (*Eristalis tenax*) over the winter, and all of the results reported herein are obtained with this species.

3.1.1 STMD Responses to Artificial Imagery

- Several recordings were obtained from STMD cells in *Eristalis* specimens, with similar responsiveness to small moving targets against moving cluttered background, as reported for the first such cell found. Analysis of data obtained during these experiments has confirmed the prior observations and revealed more detail about the properties of this neuron. It is selective for direction of target travel, and it appears to be capable of detecting (and responding exclusively to) small object motion for any combination of relative velocity and contrast that permits target discrimination by the human eye and brain. It completely rejects background motion alone. Similarity of characteristics in these different recordings, as well as the common anatomical location, suggest that a stereotypical, identifiable neuron may be involved. Because of this observation, we are making preparations for die injection in future recordings to establish the identity of the cell. We speculate that it may be a large male-specific lobula neuron which has been identified in the past, but whose response capabilities with respect to small moving targets has heretofore been unknown.

Dr. O'Carroll performed analysis of data obtained by Dr. Tamath Rainsford on this task.

3.1.2 STMD Responses to Aerial Imagery of Moving Targets

During the course of the project, we prepared stimulus data specifically tailored for the application that is the ultimate aim of this project, by importing the air-to-ground video data provided by NAWCWD, and formatting the data use in the VisionEgg, the automated stimulus generation software packaged developed in the O'Carroll laboratory. Processing of this data included frame interpolation to allow the 200 fps frame rate necessary to achieve flicker fusion in the fly eye. This enables display of these video sequences as direct stimuli for insects.

The sequence used for the bulk of the experiments was that stored in the file 1527-1.cff. This sequence consists of a truck moving along a straight section of road, filmed from a aerial camera platform that is directed towards the vehicle. Re-scaled to our stimulus display size, the vehicle's front left tire subtends an angle of approximately 1 degree at the insect eye, a not unrealistic scenario given the camera field of view. On the grounds that the truck was high in contrast and of similar size to targets that elicit strong responses from the neurons, we selected this image sequence from among several other candidates for our pilot experiments.

During the first half of the project, we presented this image sequence in a number of different recordings from candidate STMD neurons, during which we were unable to observe any significant increase in neural activity due to target passage. We suspected that this failure to respond might arise due to one or more causes. One had to do with the interpolation technique used to obtain the high frame rate (200fps) necessary to exceed the flicker fusion frequency in fly

eye: this method still allowed periodic dwell on stationary images, with the subsequent 'jump' of a moving feature from frame to frame generating transient (flicker) artifacts in neurons associated with early visual processing. The response to these artifacts could interfere with detection of a small moving feature. (If similar imagery were available at a higher original sample rate (e.g. from a high-speed video source) this would be a more realistic set of images to use in such experiments). Secondly, the images were initially presented full-frame and in a fixed location on our display system. Individual STMD neurons, however, have localized receptive fields and strong direction preferences. The imagery chosen for this experiment was selected partly because that trajectory and location on the screen match the preferences of a particular STMD neuron that we have recorded from on several occasions in the past. However, in our initial experiments we were not successful in penetrating this cell, and have recorded instead from other neurons for which the target track may not be optimal with respect to the excitatory receptive field. Finding and recording from a neuron with a suitable receptive field is a very difficult task, particularly given the short duration of viability of STMD cells following electrode penetration.

Given these issues, in the latter part of the project we developed the means to manipulate the positioning and orientation of the moving imagery on the CRT screen during experiments. This allows the adjustment of the imagery to match the receptive field location and properties of individual cells, which are probed in the initial stages of the experiment with a series of artificial small target motions in order to (at least roughly) characterize their receptive fields.

A final set of experiments performed with this flexible stimulus system yielded positive results. On two occasions, we succeeded in recording from an STMD neuron that responded unequivocally to the moving truck target. A video segment depicting the results of two trials from one of these experiments is supplied with this report, and is entitled `movie2_3_sor3_400kps.mov`. (Video segments along with a note on codecs/players are transmitted in an archive file entitled `movies_N00014-03-M-0171.zip`.)

In this video, the upper panel shows the membrane potential of the STMD neuron, which is a spiking neuron (i.e., one that generates action potentials when excited). A dense train of spikes in this panel therefore represents a strong response. The bottom panel shows the stimulus imagery. The receptive field of the neuron is in the upper center portion of the display. When the video begins, the background is pitching rightward (the preferred direction for the STMD neuron) and slightly downward, and the cell reacts strongly, apparently to the motion of the bridge shadow and dark area near it. It isn't entirely certain why there is a response, since this represents background motion, but we believe it may be a 'startup transient' at the onset of motion (which we see in simulations as well), and it is probably compounded by the fact that the shadow displays much higher contrast than the rest of the background. Following that, there is a strong response to the motion of the truck, followed by a weaker response to the small car behind it. After the car enters the receptive field, the camera pitches leftward, greatly reducing the velocity of the targets on the retina, and the cell response all but extinguishes.

These results demonstrate that the insect visual system is capable of detecting small moving targets in imagery of relevance to the air-ground scenario. Data on this task were obtained by Drs. O'Carroll and Rainsford.

3.2 STMD Modeling and Simulation

This task has seen the bulk of activity under the contract. We have developed some significant new tools for the generation and handling of insect visual data, and have run

simulations in which the central characteristic of the 'Moving Target / Moving Background' property have been modeled.

3.2.1 Insect Vision Data Tools

We have developed a standard format for storage and interchange of insect visual data, and have also developed tools built around this format, for the generation and display of data and for creation of simulation inputs or conversion of simulation outputs. These tools are primarily in the form of code written for the standard mathematics package, Mathematica.

The standard format was developed by Dr. Shoemaker in consultation with Dr. O'Carroll. Tools for processing data were developed by Dr. O'Carroll, and Dr. Thomas Bartolac of Tanner Research.

3.2.1.1 'Bug's Eye Data' (BED) Format

This is a standard format for storage and interchange of insect visual data. It specifies a convention for a hexagonal (ommatidial) grid on which retinotopic signals are defined, the data and file formats to be used for interchange, and other details of data structure. It is designed for time-sequences of data corresponding to frames in a video sequence. It is intended for any retinotopic data (i.e., raw input data, or processed data anywhere along the visual pathway in which retinotopy is maintained). A document containing the specifications is supplied with this report as Appendix A.

3.2.1.2 Data Generation / Conversion Tools

These tools sample video data from arbitrary scenes (either natural or artificial) onto a hexagonal grid (with appropriate spatial filtering), and generate BED datafiles.

3.2.1.3 Simulation I/O Handling Tools

These tools convert data from BED format (for example, raw video data) into formats required for input to simulation packages (principally SPICE, a circuit-oriented tool, and MatLab). They also convert (retinotopic) simulation outputs into the BED format.

3.2.1.4 Movie Generation

We have developed the capability to create movies from BED and simulation output data, for visualization and display purposes. This process involves generation of hexagonal metapixels and writing of frame data in the form of bitmap files using Mathematica, and the assembly of these bitmaps into a movie using Virtual Dub, another commercially available tool.

3.2.2 STMD Modeling

Evidence suggests that small target motion processing occurs in several discrete anatomical / physiological stages. The first we believe to be an elementary motion detection operation that may be common to all visual motion processing in the insect brain. At the start of this project, based on prior work for AFOSR, we postulated the existence of two additional stages dedicated specifically to small moving targets: an *elementary small target motion detector* (ESTMD), with relatively limited receptive field and retinotopic distribution, and wider-field STMD neurons that act as collators in some sense of the outputs of ESTMD cells. Experimentally, we have observed

cells whose response characteristics appear to be consistent with each class of cell under this hypothesis.

However, we do not think this hypothesis is sufficient to explain the 'moving target / moving background' capability, and we are now considering a more elaborate series of elements and operations in the processing chain that ultimately leads to a wide-field STMD neuron with the 'moving target / moving background' capability. In our simulation and modeling for this project (mindful of limited time and resources), we focused on one of these elements – the one most critical for distinguishing target from background motion, particularly when they are in the same direction.

The bulk of the STMD modeling reported herein has been at an abstract level and carried out by Dr. Shoemaker under support of the current contract. Some related, more biologically detailed modeling has been done by Dr. O'Carroll under support of US Air Force Office of Scientific Research contract F49620-01-C-0030.

3.2.2.1 *Model Structure*

Below for completeness are listed all of the elements under consideration:

1. Elementary motion detector (EMD): We have proceeded with modeling on the assumption that small target motion detectors ultimately rely for their primary inputs on elementary motion detectors (EMDs) of the correlational or Reichardt type. There is some evidence to suggest that this is the case, although it is by no means definitively proven. In particular, excitatory responses of STMD neurons are typically (although not always) strongly direction-selective, and in some cases are tuned to the velocity of target motion (i.e., give optimum response to a particular velocity, rather than responding monotonically). STMDs always give responses that depend on the contrast between the target and background. These are all expected properties that arise from Reichardt-type motion detectors.
2. Longitudinal target size discrimination: Ability to discriminate 'smallness' of target in direction of motion. A poster by the collaborators on this project [5] gives some indirect evidence for processing of this nature in the responses of STMDs, and reports an initial effort to model it. Such processing may serve to 'prefilter' the outputs of EMDs for events consistent with passage of a small object, while rejecting the effects of moving regions with greater length. Because these concepts are still in development, this element has not been incorporated in simulations done for this project.
3. Target / background speed discrimination: An essential capability for distinguishing a moving target from a moving cluttered background. This element was a major focus of work under this project.
4. Lateral target size discrimination: Ability to discriminate 'smallness' of target perpendicular to direction of motion. This (along with enhancement of directional selectivity) may be a principal function of neurons that we have previously labeled elementary small target movement detectors (ESTMDs). They display a selectivity for laterally 'small' targets by virtue of their receptive field shape in visual space: a 'notched' inhibitory region surrounds an excitatory receptive field, so that anything that doesn't fit through the notch causes inhibition. This element, although briefly considered, has not been incorporated in simulations done for this project.

5. 'Dendritic processing': A mechanism in a wide-field collator cell for stimulus summation and artifact rejection. Reinforcement of either active or diffusive potentials along a dendrite is assumed to occur by sequential excitement of ESTMDs along continuous tracks in visual space. This feature exploits the constraint of continuity of motion to obtain higher-confidence detection of a moving target. Processing of this kind was considered as an essential stage following target background speed discrimination, although it was not explicitly simulated.

3.2.2.2 Target / Background Speed Discrimination

This particular element of our STMD modeling received the major focus of attention under this effort. Early on in the project, we considered two hypotheses with respect to the capability of discrimination on the basis of *velocity*, between motion of local features in the background (which can look a lot like small targets), and motion of a small target that is moving with respect to background on the retina:

1. STMD cells 'look for small-target events' in the outputs of *individual* EMDs, and judge whether these events are consistent with the velocity of local background motion.

2. STMD cells 'look at' spatiotemporal sequences of 'small target events' in two or more adjacent EMDs, and judge whether this sequence is consistent with the velocity of local background motion. Several possible ways to compute 'Consistency with velocity of local motion' are under consideration, and are discussed in more detail below.

Subsequent work has led us to a model that is based on the second of the two hypotheses above. (We found no evidence to suggest that the output of an individual EMD carries sufficient information to discriminate the passage of a small target against a moving background, under general conditions.) Our current approach involves a comparison of the traversal times across pairs of ommatidia of 'target events' with the expected traversal times for moving background (which might well have features like small targets). This comparison is achieved with an EMD-like element that is tunable by feedback representing state of local motion, and which performs anticorrelations rather than correlations to look for events inconsistent with that motion. This element is referred to as a 'primitive STMD', or PSTMD.

The operation of the PSTMD model is best explained by first considering a one-dimensional or uniaxial case. Target and background are assumed to move at different speeds along this axis. Input signals from early vision are first passed through elementary motion detectors (aligned in the same direction), which give enhanced response to moving stimuli while rejecting flicker or non-motion-related temporal contrast. The EMD also eliminates dependence on contrast polarity – the passage of bright targets or edges gives the same response as dark targets or edges. Instead, the polarity of the EMD output is dependent on *direction* of motion: the strong initial transient caused by event passage is positive in the 'preferred' direction of the EMD and negative in the 'antipreferred' direction. This antisymmetric, bipolar response allows an EMD to carry information about motion in either direction with respect to its axis of alignment.

Thus, when the target and background are traveling in opposite directions, the target event is relatively easy to pick out in the 'motion imagery' of an array of EMD outputs: its polarity is opposite to the dominant image polarity. When the target and background move in the same direction, however, a mechanism to explicitly differentiate between them is necessary. This is the function of the PSTMD.

In spite of its advantageous features, the EMD does have some disadvantages with respect to its function as the front-end in small target motion detection. One is the dependence of its response with respect to target speed: an EMD has an 'optimum velocity' that elicits the strongest response, and if background is moving near this velocity while the target is moving at some sub-optimal velocity, the subsequent stages will be presented with a higher degree of background 'noise' from which the relatively weak target response must be distinguished.

In our modeling to date, the PSTMD operates locally on the outputs of pairs of EMDs. (We expect, however, that the function of longitudinal size discrimination may actually be interposed between the EMD and the PSTMD, and this is likely to significantly improve its discrimination for small targets and reduce artifactual responses.) A PSTMD delays the input from an EMD that is 'upstream' with respect to its preferred direction, and uses this delayed signal to inhibit its own response to the second, 'downstream' EMD input. (Note that, unlike the EMD, the PSTMD does not have an inherent antisymmetry, and is thus configured to detect motion in only one direction along its axis.) The delay is tuned by some adaptive mechanism so that it represents the expected delay due to the local background motion. We have considered subtractive and divisive mechanisms for inhibition, and have focused mainly on division, which models a biological shunting inhibitory mechanism. PSTMD state is computed as the quotient of the downstream input by the delayed (absolute value) upstream input plus a small constant, which represents the inverse of the maximum gain for the PSTMD. The 'delay' operator is not a pure time delay (although if it were, the response of a properly tuned PSTMD would be constant for uniform background motion); rather it is a continuous-time operator corresponding to a lowpass filter, which is more readily believable with respect to the biology, and more amenable to asynchronous analog implementation in integrated circuitry. In our case, we have used a second-order linear filter with a complex pole-pair.

This PSTMD model responds to uniform background motion with relatively minimal variations in its output, but when events at the upstream and downstream inputs do not 'match up', as is the case when a target is passing with different speed than background, a more significant response is generated. Naturally, of course, the target must have some contrast with the local background at the time of its passage; no algorithm is capable of detecting a target that cannot be distinguished from background.

The PSTMD concept is fairly straightforward in a one-dimensional scenario, but is less so for two dimensional imagery and sensing. In the real animal, there are multiaxial EMDs, and any background motion will necessarily be poorly aligned with at least some of the interommatidial axes. Consider the response of a PSTMD as the direction of local background motion varies with respect to its axis in retinotopic space. When the two are aligned, the analysis above for the uniaxial case applies. As the angle between the direction of motion and the axis increases, the degree of correlation between events at the two PSTMD inputs decreases due to aperture effects, and the time delay between events which are correlated decreases as the cosine of the angle. When the motion direction is at right angles to the axis, then only objects as wide as the inter-input distance will cause correlations, and the average time difference between such events will be zero. (Obviously, if the inputs to the PSTMD are prefiltered for 'small target events', then response to such larger objects may be excluded entirely.) For angles greater than 90 degrees in magnitude, we do not assume that the systems 'tunes' with 'negative' time delays, but rather that the difference in signs between small target and background events can be exploited to pick out

the target, and the time constant may assume some default value with little effect on the target detection process.

Because the degree of correlation in two inputs to the PSTMD decreases as the angle between its axis and the background motion increases, there will be less perfect inhibition of the PSTMD response by its delay-and-inhibit mechanism, and the incidence of large but spurious responses that could be mistaken for small target events increases. Thus there is a need for higher-order correlation of such events; a model for such a function in a collator neuron is discussed in the next section.

How such a PSTMD might be tuned is still an open question. We are considering three possible mechanisms (although none was sufficiently developed to evaluate in simulations during the short duration of this contract): one involves computation of a signal based on the outputs of EMDs in the local area; the second, some form of local feedback approach that seeks to minimize the outputs of the 'anticorrelating' PSTMDs themselves over time, and a third uses a feedback signal derived from a set of wide-field 'tangential cells', a class of neurons that in insects are presumed to carry information about the global state of egomotion from which the local expected velocity could be derived.

Finally, the separation in retinotopic space of the 'upstream' and 'downstream' inputs to a PSTMD is an issue that will require some further study and analysis. Larger separations permit target discrimination for smaller relative velocities between target and background, but also result in sharper direction-tuning (i.e., selectivity for motion in a particular direction). The latter characteristic needs to be matched to the angle between different PSTMD axes (which we take in our models to be the hexagonal angle, or 60°). In our modeling, we have assumed that the input separation is equivalent to two ommatidia, an *ad-hoc* choice that seems to work satisfactorily.

The response of a PSTMD unit is not dependent on the width or lateral extent of a moving object that excites it: a traveling bar would excite entire rows or columns of PSTMDs. Thus, prior to collation by a wide-field STMD neuron, a processing stage may be present in the STMD chain that rejects broad moving targets, as discussed in Section 3.2.2.1 above.

3.2.2.3 Dendritic Processing

Now consider the responses of a series of PSTMDs that lie along the track of a small target as its image moves across the retina. When the background motion is in the same direction, its effects are repressed, and the PSTMDs will respond strongly only to the target, with a series of activations along the target track. When there is significant misalignment between the background and target motion, non-correlated events will occur due to the background motion and will lead to spurious transients that may be indistinguishable from the response to a target. However, such events will follow a track in retinal space that is in the same direction as the background motion and is misaligned with the track of target passage. The greater the misalignment of target and background motion, the more spurious events will be generated, but the more their track will deviate from the target direction.

We consider as the final stage in STMD processing a collating neuron (most STMD neurons studied experimentally have moderate to large receptive fields in visual space, suggesting that such cells are collators of local responses), and we hypothesize that a cell of this type performs spatiotemporal summation of PSTMD outputs in a way that reinforces its response when those outputs occur along a continuous or near-continuous track in the preferred direction. This could occur by an ordered pattern of synapses from the PSTMDs onto dendrites in the target neuron,

such that reinforcement of either active or diffusive potentials along the dendrite occurs by sequential excitement of PSTMDs on a preferred-direction track in visual space. This would serve to reject artifactual events, and would also exploit the constraint of continuity of target motion (or near continuity, if small obstructions or equiluminant background objects were present) to obtain higher-confidence detection of a moving target. The cell output would presumably be formulated by thresholding the internal cell state, and the cell would only reach activation when multiple, sequenced events occur on preferred-direction tracks.

The spatiotemporal properties of the collator dendrites affect the tuning of the cell to speed of target motion, but we assume such tuning to be quite broad. This can be accomplished, for example, by active membrane properties that result in fast onset and slow decay of excitatory potentials. In this way, the cell responds to any series of synaptic events whose individual duration is sufficient to elicit the rapid activation, but whose timing is faster the much-longer decay. Evoked potentials with fast onsets and slow decay have been observed in the neurons we have labeled ESTMDs, and we have modeled them abstractly and in analog silicon. Detailed results will not be given herein, but responsiveness to a tenfold range of target speeds is readily practical.

3.2.3 PSTMD Simulation

An array of PSTMDs, preceded by processing that mimics early vision and elementary motion detection (the correlational EMD), was simulated on the standard 31 x 31 hexagonal grid mentioned in Section 3.2.1 above. In this grid, the interommatidial spacing (measured in degrees of visual angle subtended) is 1.73°. A triple of elementary motion detectors was defined at each hex pixel, aligned with the interommatidial axes. PSTMDs with preferred directions in three of the six possible interommatidial directions (downward, upper right to lower left, and lower right to upper left) were also defined at each hex pixel. Simulations were run on both natural and artificial imagery. Artificial imagery was produced with the use of the "VisionEgg" stimulus generation software developed in the O'Carroll lab at Adelaide, and processed for blurring (to mimic compound eye optics) and hexagonal resampling. In addition, video imagery supplied by NAWC was processed in the same manner. Simulations were performed at Tanner Research using the tool SPICE (Simulation Program with Integrated Circuit Emphasis). SPICE is based on circuit representations, but admits the use of purely abstract components as well as particular electronic devices. Such abstract components were used to implement all processing elements for purposes of simulating STMDs.

These simulations were carried out by Dr. Shoemaker, with Dr. O'Carroll providing moving image data for inputs to the simulated arrays.

3.2.3.1 Simulations with Artificial Imagery

In these simulations, the stimulus consisted of a small dark target (just over one hex pixel in size) orbiting clockwise in a circle about 15 ommatidia in diameter at a speed of about 40°/s, where degrees refers to visual angle subtended. (This speed is a bit under half the optimum speed for excitation of the front-end EMDs). This target motion has the advantage of covering all directions of travel in a single data segment. It took place against uniform translatory background motion of various speeds and directions. The animated background was a (spatially broadband) random textel pattern with statistics similar to natural scenes, with an average contrast of 40%.

In these simulations, the tuning of the PSTMD elements was done by hand. The delay time constant for a PSTMD aligned with any particular axis was set in proportion to the spatial basis of the PSTMD (two ommatidia), divided by the background speed and the cosine of the angle between the direction of motion and the axis. When that angle was very close to or greater than 90° , the time constant was set to a small positive value. The Q factor for all PSTMD delay filters was set to unity.

These results have confirmed the properties of the PSTMD deduced by the qualitative analysis of Section 3.2.2. In particular, the PSTMD is successful in suppressing background motion and accentuating a target moving at a different speed, when both are in alignment with the PSTMD axis. In Figure 1 below is shown the output (as a function of time) of a downward-sensitive PSTMD which is located at the '3-o'clock' position on the orbit of the small target. In this case, the background speed is about twice the target speed. In the trace, a pronounced transient response due to the passage of the target moving downward through the three-o'clock position on each of two orbits is visible.




Figure 1: Response of downward-sensitive PSTMD unit to moving target passing downward at times 0.5s and 2.5s against fast downward-moving background.

In Figure 2 below, another PSTMD output from the same simulation is depicted. This PSTMD is aligned for upper-right-to-lower-left target motion (60° orientation), and located at the '5-o'clock' position on the orbit of the small target, when the target is moving in its preferred direction. The target passage occurs at times 0.83s and 2.83s. Note that larger transients due to background motion are present in this output as compared to that in Figure 1, due to poorer correlation of events in the delayed and undelayed inputs to the PSTMD.




Figure 2: Response of PSTMD unit tuned to upper-right-to-lower-left target motion, with a small target passing in that direction at times 0.83s and 2.83s. Background motion is as in Figure 1.

In Figure 3 below, a third PSTMD output from the same simulation is depicted. This PSTMD is aligned for lower-right-to-upper-left target motion (120° orientation), and located at the '7-o'clock' position on the orbit of the small target, when the target is moving in its preferred direction. The target passage occurs at times 1.17s and 3.17s. Note in particular that the artifactual responses due to background motion are predominantly negative, because the alignment of the PSTMD varies from the direction of background motion by more than 90° . This makes the positive responses to the target easier to discriminate.




Figure 3: Response of PSTMD unit tuned to lower-right-to-upper-left target motion, with a small target passing in that direction at times 1.17s and 3.17s. Background motion is as in Figure 1.

The most difficult scenario for the PSTMD model (that from which the results above were taken) proved to be the case when in which target and background move in the same direction, and the background motion is faster and closer to the optimum velocity of the front-end EMDs than is the velocity of the target. In this case, the artifacts present in the PSTMD outputs are relatively large, compounding the problem of occasional obscuration of the target by equiluminant background. Higher-order correlators would definitely be required in such a case to pick the target motion from the background in the PSTMD outputs.

A video segment depicting the results of another simulation is supplied with this report, and is entitled `pstmd_slow_data1c (DivX5.1 25fps).avi`. These results are taken from the case when the background is moving downward, but with a slower speed than the target. The video display consists of four panels depicting, from left to right, the 'insect eye' view of the raw scenery (i.e., blurred and hexagonally resampled), and the outputs of PSTMDs aligned with the downward,

upper-right-to-lower-left, and lower-right-to-upper-left directions, respectively. In the PSTMD false-color imagery, 'warm' colors indicate positive output values, and 'cool' colors indicate negative output values. The strong peaks in the response due to target passage in or close to the preferred direction of each PSTMD are evident (although occasionally weakened by target obscuration by dark portions of the background).

3.2.3.2 Simulations with Natural Imagery

An STMD simulation was run using a segment of an air-to-ground video supplied by NAWC as the input data, taken from an air platform, and depicting a highway along which a (relatively large) truck target moved. Due to non-uniform background motion in this data, appropriate tuning could not be maintained for all PSTMDs, and the tuning was roughly estimated for background motion from upper right to lower left, which occurred during initial appearance of the target. Nonetheless, the target is clearly highlighted in the PSTMD imagery (although other high contrast features in the imagery cause responses also). A video segment depicting the results of this particular simulation is supplied with this report, and is entitled `pstmd_array_global_T_data3a (DivX5.1 25fps).avi`. The format is the same as for the video described in Section 3.2.3.1.

3.3 Higher-Level Issues

3.3.1 Integration of Model Elements for STMDs

In Section 3.2.2.1, we discussed elements of an STMD model in terms of particular functional characteristics of those elements. We hypothesize that these might be integrated into a complete, useful STMD as follows. The order in which the functional elements were presented in that section is assumed to represent the processing order performed by a series of physical processing stages (which are quite possibly anatomically distinct neurons or neuron types in the biological system). The outputs of elementary motion detectors would first be processed by filters that are tuned particularly to the relatively short-duration events that would be consistent with the passage of a small target. Naturally, such events could also be present due to background motion, as caused by the presence of small, fixed objects in that background. However, by eliminating responses to larger objects, the inputs presented to subsequent processing would induce fewer artifacts. The next stage in the processing would be a PSTMD array, followed by an array of filters which would eliminate responses to laterally extensive targets in the PSTMD imagery. The reason that this processing step would follow the PSTMDs is that the candidate neuron for this function, the ESTMD cell, uses a mechanism that would result in strong inhibition if it operated on imagery in which background motion events were extensively present. Finally, the PSTMD cell outputs would be collated by wide-field STMD neurons with directional dendritic processing as described in Section 3.2.2.3.

Implementation of these elements in an artificial system could be accomplished most naturally (and with the smallest volume and power consumption) by vertical integration of semiconductor chips, as discussed below in Section 3.3.2.2.

3.3.2 Physical Implementation of Artificial STMDs

With respect to implementation technologies for artificial STMDs, we consider asynchronous, neuromorphic analog VLSI to be a medium of choice for direct implementation. Size, power, and

eventual cost advantages would accrue. In ongoing research (US Air Force SBIR contract F08630-02-C-0013), we are developing a VLSI implementation of an array of adaptive EMDs, which we expect would serve as a front end for STMD processing.

3.3.2.1 STMD Functions in Analog VLSI

Under support of our second Air Force contract F49620-01-C-0030, we have developed circuitry to emulate the processing performed by the neurons we have labeled ESTMDs, and to compute the spatiotemporal correlations that were discussed in Section 3.2.2.3 with respect to dendritic processing in collator cells. The silicon ESTMD model involves an isotropic, diffusive inhibitory network, which connects onto excitatory 'neurons' in an anisotropic manner to implement directional inhibition. This 'neuron' also receives direct excitatory input from one or a few adjacent EMDs; the combination of excitation and inhibition leads to a receptive field with an excitatory center and a 'notched' inhibitory surround. The dendritic model allows potentiation of waves of excitatory input by sequential activation of input circuits that receive inputs from spatially adjacent sources, and that feed forward to subsequent circuits in the 'dendrite.'

A key component of both models is a nonlinear temporal filter circuit that responds to (unipolar) inputs with fast onset and slow decay behavior.

Under a renewal project proposed to the Air Force Office of Scientific Research, we will begin modeling the other STMD functions discussed in Section 3.2.2, in analog silicon.

3.3.2.2 Layered Processing / DARPA VISA Program

With a low-power analog VLSI approach to neuromorphic circuit design, synchronous, high-bandwidth interconnection between chips (i.e., readout and read-in circuitry) creates serious challenges with respect to contamination of analog signals with digital noise, and increases power consumption. Direct, parallel interconnection is preferable, but this approach requires vertical integration of semiconductor chips. In such an approach, chips are thinned, vias are etched through and lined with insulation (most likely grown or deposited oxide), metal is introduced to fill the vias and contact I/O nodes in the circuitry on chip, and then finally, chips are aligned and bump-bonded together. Vertical integration technology is being studied and developed under a DARPA/MTO program entitled VISA (vertical integration of sensor arrays). This program addresses parallel interconnection of high-density sensor arrays and readout circuitry, but could equally be applied to analog neuromorphic processors in multi-chip systems. Because the desired interconnections in the sensor array problem are dense, the technology may actually find easier and shorter-term application to insect-vision-based artificial systems, which at this point in time would have a relatively low interconnect density.

3.3.3 Application of STMD Technology to the Air-Ground Seeker Problem

The characteristics of STMDs suggest that they may play a role in multiple phases of target pursuit, including acquisition, intercept, and closing pursuit. Some STMDs have large receptive fields and are exquisitely tuned to small target motion; such cells may play a role in simply 'sounding the alert' on the presence of a target. Others with smaller receptive fields may serve to localize moving targets during flight. Once acquired, evidence suggests that dragonflies use a form of proportional navigation during pursuit. Many other insects, however, seem to acquire as accurate a fix as possible on target speed and direction, and compute an intercept that is flown in

large part in an open-loop fashion. Below are discussed in further detail some considerations for the role of the STMD in pursuit behavior.

3.3.3.1 *The Insect Vision Paradigm and Flight Control*

A major assumption that is a basis for the current effort is that STMD cells of some type play a major role in target tracking and 'seeker' functions as well as detection. This is suggested both by their fascinating response characteristics and the fact that they appear most prevalent in insects that chase prey, mates, or rivals on the wing. However, this hypothesis has significant implications with respect to control of the behaviors involved in tracking small targets in any closed-loop fashion. While details of control theory are beyond the scope of this Phase I effort, some considerations of the implications of the sensory processing are in order.

Conventional tracking and seeker techniques typically operate globally on an image or subimage. When biomimetic or other local motion processing is used (e.g., the EMD or other motion-energy-based formulations, or the Wave Process algorithm [1] developed at Tanner Research), the goal is typically image enhancement to highlight target motion, or identification of regions of interest to reduce data throughput. Subsequent operations rely on conventional target state estimation fed to a controller. However, any STMD-based system in the insect must differ from this paradigm: when regarded in the context of target state estimation, it necessarily involves a rather coarse quantization of state variables associated with the fact that target estimation is coded by a finite ensemble of discrete neurons. Each of these 'looks at' a limited receptive field, and the response of each is built up from highly local calculations (in the EMDs). Target location information is place-coded; that is, indicated by the particular cells activated, since the receptive fields of STMDs vary about the visual sphere from cell to cell. The spatial resolution of such coding is presumably much lower than that of the (already low) retina. Direction of target motion is also place-coded by at least some of the STMD cells, since many of those we have examined are directionally selective. There is no evidence as yet that direct information about target state (such as velocity) is carried in the analog responses of the individual STMD neurons themselves (although responses of individual neurons are strongly tuned to velocity, so that 'higher order' comparisons could deduce this variable); an excitatory response in an individual may simply indicate the presence of a target moving in an area and direction, and within a range of speeds, for which the cell is tuned. It is, however, possible that the response strength is related to some parameter(s) of estimation rather than parameters of motion itself: we have observed that responses in some STMDs are stronger when the target motion is unambiguously visible against background, and weaker when the target is more difficult to pick out (by the human eye), due to similarity to background features and/or low relative velocities between target and background.

Overall, in its position-coding, this scheme is more akin to a simple four-quadrant sensor / seeker (left/right in azimuth and up/down in range) than more sophisticated artificial seeker algorithms, although the quantization with respect to target position is probably somewhat finer. The basic response characteristic differs from the position-based sensor / seeker in that it is velocity dependent: target motion is *necessary* to induce any output at all. The system must therefore rely on slippage of a pursued target relative to both background and in absolute terms on the retina, and it cannot simply "servo" target position. This would seem to imply that some sort of active vision is necessary during pursuit, at least up to the point at which an intercept may

be estimated and executed in open-loop fashion (which does seem to occur in the terminal phases of some insect pursuits).

3.3.3.2 Alternate Paradigm: Primitive STMD as Front-End for Conventional Processing

An alternate to a fully biomimetic approach is to use the PSTMD as a preprocessor for conventional target state estimation or other seeker functions. In this context, it would serve a small-target enhancement function, much as do elementary motion detection or other motion-energy-based formulations, or the Wave Process algorithm in current scenarios. Such an approach offers the possibility of near-term application that may be practical in the short-term duration of a Phase II project.

3.3.3.3 Specific Issues for Air-Ground Platform (LOGIR)

With the LOGIR platform, features that are intended to reduce cost and complexity may be (unintentionally) problematic to our biomimetic approach. These features include the use of non-gimbaled sensors (there is good evidence that insects stabilize regions of their visual field while in flight, and saccade between such stable gazes; this would be an advantageous strategy in a biomimetic system), and the use of commercial uncooled IR sensors (which require readout and read-in circuitry). Ultimately, parallel integration of sensor and the first stages of biomimetic processing would be a much more advantageous approach. (see Section 3.3.2).

4 Follow-On Work

In our work, we have developed simulation, modeling, and data-handling tools essential to continued investigation of a biomimetic approach to the problem of a low-cost, optically-based seeker for autonomous weapons and platforms. More importantly, we have made a first effort at modeling an essential computational component in the chain of motion processing that leads to the 'moving target / moving background' capability, which lays the groundwork for integrated models of this processing that can be applied to artificial systems. This work represents an important first step (albeit an early one) in learning how to apply biological models for moving target detection to biomimetic solutions to real-world problems. We believe any follow-on work should be thoroughly integrated with a broader research and development program that focuses on both biomimetic processing and specialized hardware development, and which is also being supported or will be supported by complementary research contracts. This coordination will be necessary to transition the technology into applications, and bring it to a level of maturity that will allow its evaluation for real weapons systems.

The next phase suggested by this work is an effort that is focused on several considerations: 1) the continued development of implementable STMD models; 2) the development of specialized hardware for STMD processing; 3) the issue of integration with sensing and early visual and motion processing (the EMD); and 4) integration of small moving target detection with the higher-level processing / decision-making and flight control systems that will be necessary to apply the STMD to the problem at hand. These research elements would leverage heavily off of other ongoing or proposed work: development of implementable STMD models will take place in conjunction with parallel research on the biological systems (conducted under support of AFOSR contract F49620-01-C-0030 and follow-ons); the development of STMD

hardware will be based largely based on silicon models developed under support of the same Air Force program; and, integration with sensing and the EMD will leverage off of other Air Force-supported efforts (SBIR contract F08630-02-C-0013 and follow-ons).

Bibliography

- [1] Bartolac, T. "Distributed predetection of moving point targets with analog VLSI," *SPIE, Infrared Readout Electronics*, Vol. 1684, pp. 267-281, 1992.
- [2] Buchner, E. "Elementary movement detectors in an insect visual system," *Biol. Cybernetics* Vol. 24, pp. 85-101, 1976.
- [3] Collett, T.S. and Land, M.F. "Visual control of flight behaviour in the hoverfly, *Syrirta pipiens*," *J. Comp. Physiol. A*. Vol. 99: pp. 1-66.
- [4] Douglass, J.K., and Strausfeld, N.J. "Early motion processing pathways in insects," In Zanker, J.M. and Zeil (eds), *Computational, neural and ecological constraints on visual motion processing*. Springer Verlag, Berlin, 1999.
- [5] DuBois, R.A., O'Carroll, D., and Shoemaker, P.A., "Spatio-temporal tuning for small targets from a simulated array of elementary motion detectors," in N. Elsner & H. Zimmermann, Eds., *The Neurosciences - From Basic Research to Therapy: Proceedings of the 29th Göttingen Neurobiology Conference* (Germany), Georg Thieme Verlag, Stuttgart, pp. 547-548, 2003.
- [6] Egelhaaf, M., Borst, A. & Reichardt, W. "Computational structure of a biological motion-detection system as revealed by local detector analysis in the fly's nervous system," *Journal Of the Optical Society Of America A - Optics Image Science and Vision* Vol. 6, pp. 1070-1087, 1989.
- [7] Egelhaaf, M. "On the neuronal basis of figure-ground discrimination by relative motion in the visual system of the fly II. Figure-detection cells, a new class of visual interneurons," *Biol. Cybernetics* Vol. 52, pp. 195-209, 1985.
- [8] Egelhaaf, M., and Borst, A. "A look into the cockpit of the fly: visual orientation, algorithms, and identified neurons," *Journal of Neuroscience* Vol. 13, no. 11, pp. 4563-4574, 1993.
- [9] Harris, R.A., O'Carroll, D.C. & Laughlin, S.B. "Contrast gain reduction in fly motion adaptation," *Neuron* Vol. 28, pp. 595-606, Nov. 2000.
- [10] Harris, R. A., O'Carroll, D. C. & Laughlin, S. B. (1999). "Adaptation and the temporal delay filter of fly motion detectors," *Vision Research* Vol. 39, 2603-2613
- [11] Hassenstein, B., and Reichardt, W. "Systemtheoretische analyse der Zeit-, Reihenfolgen-, und Vorseichenauswertung bei der Bewegungspereption des Rüsselkäfers *Chlorophanus*," *Z. Naturforsch.* Vol. 11b, pp. 513-524, 1956.
- [12] van Hateren, H. "Directional tuning curves, elementary movement detectors, and the estimation of the direction of visual movement," *Vision Research* Vol. 30, No. 4, pp.603-614, 1990
- [13] Henry, J. & O'Carroll, D. "Visual detection of prey by dragonflies," Proceedings, International Conference on Invertebrate Vision, Bäckaskog Castle, Sweden, 7-12 August, 2001.
- [14] Krapp, H. G., Hengstenberg, B. & Hengstenberg, R. Dendritic structure and receptive-field organization of optic flow processing interneurons in the fly. *Journal of Neurophysiology* Vol. 79, 1902-1917, 1998.
- [15] Krapp, H. G. & Hengstenberg, R. Estimation of self-motion by optic flow processing in single visual interneurons. *Nature* Vol. 384, 463-466, 1996.
- [16] Mead, C. *Analog VLSI and Neural Systems*, Addison-Wesley, Reading, MA, 1989.
- [17] O'Carroll, D. & Laughlin, S. "Asymmetric receptive field organisation of neurons in the insect lobula". Proceedings of the 24th Göttingen Neurobiology Conference, N. Elsner & H-U. Schnitzler, Eds., Georg Thieme Verlag, Stuttgart & New York. 1996, Vol. II, p. 341, 1996

- [18] O'Carroll, D.C. "Retinotopically organised cells in the fly lobula with small receptive fields and 'high order' properties." Proceedings of the 22nd Göttingen Neurobiology Conference 1994, N. Elsner & H. Breer. Eds., Georg Thieme Verlag, Stuttgart & New York. Vol. II, p. 454, 1994.
- [19] O'Carroll, D. "Feature-detecting neurons in dragonflies," *Nature* Vol. 362, pp. 541-543, 1993.
- [20] Olberg, R.M., Worthington, A.H. & Venator, K.R. "Prey pursuit and interception in dragonflies". *J. Comp. Physiol. A* Vol. 186, pp. 155-162, 2000.
- [21] Potters, M., and Bialek, W. "Statistical mechanics and visual signal processing" *Journal de Physique I* Vol. 4, pp. 1755-1775, 1994.
- [22] Reichardt, W., Egelhaaf, M., and Guo, A. "Processing of figure and background motion in the visual system of the fly," *Biol. Cybernetics* Vol. 61, pp. 327-345, 1989.
- [23] Shoemaker, P.A., O'Carroll, D.C., and Straw, A.D. "Implementation of visual motion detection with contrast adaptation," *Proceedings of SPIE, Electronics and Structures for MEMS*, Vol. 4591, pp. 316-327, Adelaide, December 2001.
- [24] Warzecha, A.K., Egelhaaf, M., and Borst, A. "Neural circuit tuning fly visual interneurons to motion of small objects. 1. Dissection of the circuit by pharmacological and photoinactivation techniques." *J. Neurophysiol.* Vol. 69, pp. 329-339, 1993.
- [25] Warzecha, A.K., Borst, A., and Egelhaaf, M. "Photoablation of single neurons in the fly visual system reveals neural circuit for the detection of small moving objects" *Neurosci. Lett.* Vol. 141, pp. 119-122, 1992.

Appendix A

'Bug's Eye' Data Format, v1.0:

A standard format for interchange of insect visual data for purposes of manipulation and display, in which:

- a) Spatially discrete, retinotopic data at an instant in time are represented in a packed file of binary numbers in IEEE floating-point format;
- b) The data are associated with a regular hexagonal 'pixel' (i.e., ommatidial) grid in retinotopic space, in which location is indexed by row and column position per the convention depicted in Figure 4 below. In this convention, the pixels corresponding to the even-numbered columns are offset downward by half of the interpixel distance, relative to pixels in odd-numbered columns. The row and column indices increase from top to bottom and left to right, respectively. The default array size is 31 rows x 31 columns, although other sizes are acceptable when properly indicated per item g) below.




Figure 4: Row / column indexing scheme for data from hexagonal retinotopic grid. Columns run in the vertical or longitudinal direction, and rows in the horizontal or latitudinal direction, in the same manner as ommatidia in the frontal region of the fly eye.

- c) The order of data in a binary file is sequential and row-major with respect to the associated pixel array; i.e., from the start of the file, in the order:

Data(Row_1,Col_1) Data(Row_1,Col_2) Data(Row_1,Col_3) ...

Data(Row_2,Col_1) Data(Row_2,Col_2) Data(Row_2,Col_3)

- d) Each discrete datafile corresponds to one (and only one) frame of data, with a single variable or degree of freedom associated with each pixel. A set of such files is used to

represent a discrete-time sequence of data. By convention, the order in sequence is indicated by a number imbedded in the filename. Filenames consist of a set of (non-numeric) characters identifying and common to the entire sequence, followed by frame number indicating temporal order. The number of digits in the frame number field is consistent for all files in a particular set (i.e., leading zeros are used to fill in this field for smaller numbers), although it may vary from set to set as required.

- e) Sequences of multi-dimensional pixel data are contained in multiple associated sets of datafiles, where each set consists of a series of frames for one variable or component of the multi-dimensional data. 'Multidimensional pixel data' is used to refer to quantities, such as vectors or color components, that are amenable to representation in a unified image. By convention, the non-numeric characters in the filenames of all such associated sets of files are identical, except that the final character before the first numeric character of the frame number varies to identify the particular component or variable contained in that sequence. All such associated sets contain the same number of frames.
- f) The file extension shall be ".bed", as in bug's eye data.
- g) A set or multiple sets of data files associated a particular sequence of data is accompanied by a text file, unless the structure and parameters associated with the data conform exactly to a set of defaults specified below. The name of the text file is identical to the datafile base name, meaning the non-numeric portion of the datafile name field, excepting any final character acting as a component/variable identifier in the case of multidimensional data. The file extension is .txt. This text file is in single-column ASCII format, with each line ended by a carriage return. It contains the following:

100 lines for parameter specification (of which only the first 7 are used in this version of the format); each line contains one number:

```
#1  % .BED format version number (default=1)
#2  % byte order: 1=little-endian, 0=big-endian (default=1)
#3  % bytes per pixel, (default=4, i.e. single precision float)
#4  % number of pixel rows (default=31)
#5  % number of pixel columns (default=31)
#6  % frames per second (default=1000)
#7  % delta_phi_h: horizontal spacing of pixel rows, in terms of
    viewing angle in degrees (default=1.5; note interpixel
    separation = delta_phi_h/sin(60°))
#8  % unused
to
#100 % unused;
```

An additional (unspecified) number of lines is used to describe features or parameters particular to the data set or its applications. Users should provide a cursory description of the data sequence, and an explanation or key to characters used in the datafile names to identify components of multidimensional data.

- h) Data may be generated for multiple variables which are associated and synchronized. An example is a set of outputs of (retinotopic) elements at various points in the visual processing chain, in response to a particular moving image. This is distinct from the

concept of the multidimensional variable discussed in item e) above, which pertains to quantities intended for representation in a single image. BED datafiles for distinct but synchronized multiple variables will be given distinct names (but preferably sharing some common portion of their base names), and placed in subfolders in a directory structure for purposes of clarity. Such associated sets may share a common text file.

- i) A set or multiple sets of data files associated with a particular sequence of data may optionally be accompanied by additional files containing information pertinent to the data or its display. No specification is made at this time regarding these files, but their contents and purposes should be made clear. At minimum, an explanatory note will be included in the text file accompanying the datafiles. An example of such a file is an ASCII file specifying a color look-up table, and containing three columns of tab-delimited data representing RGB (red, green, blue) values in the range 0 – 1.
- j) Unless otherwise indicated, BED data will be interchanged in the form of zipped archives, allowing both compression of the data and maintenance of a directory structure.